

Grid Computing Center Summer Load Shed Plan

Michael Diesburg
CD Availability Manager
All Experimenters Meeting
May 6, 2013

Background

- ◆ Over the last few summers we have experienced cooling problems in the computing rooms at the Grid Computing Center (GCC) on unusually warm days.
- ◆ These issues have resulted in the need to temporarily reduce computing capacity to lighten the cooling load on hot days.
- ◆ Last summer we experienced 17 days of reduced (~30%) capacity due to these problems.
- ◆ A significant amount of effort has gone into understanding the source of the cooling load problems and planning solutions to alleviate the problem in the future.
- ◆ Basic problem is that AC condensers are on the ground behind the building. This building partially blocks air flow to the condensers. A berm behind condensers also impeded flow. In addition, under some wind conditions air patterns on the ground result in warm air being recirculated over the condensers.

GCC Load Shed Plan

As summer approaches it is likely we will again experience cooling issues at GCC

Several steps have been taken to alleviate impact of cooling problems:

- ◆ Soaker hoses used on concrete slab under condensers to lower temperatures..
- ◆ Berms behind the condensers have been removed to facilitate air flow over the condensers.
- ◆ Equipment has been redistributed in GCC to minimize impact of outages on computing capacity.
- ◆ Load shed plan has been made to handle this situation in a controlled manner that will allow partial shutdown in a planned sequence that will minimize loss of computing capacity.

Decision Process

- ◆ Facilities and Operations Department (FOD) in CD will make decision to reduce power consumption in GCC.
- ◆ Decision is at their discretion, but it is expected they will consult with wider audience if outage is likely to extend into non-working days.
- ◆ Decision will take into account following factors:
 - ◆ Time of day and forecast for temp and insolation for remainder of day.
 - ◆ Forecast for temp and insolation for non-workdays as of preceeding work day.
 - ◆ Temperature of concrete pads under the condensers.
 - ◆ Temperatures inside computer rooms
 - ◆ Any other relevant info pertinent to power and cooling
- ◆ It is expected that prudent judgment will be exercised in making this decision.

Possible Decisions

- ◆ Situation may result in one of three possible actions by FOD:
 - ◆ Do nothing.
 - ◆ Give warning of possible load shed later in the day or on coming non-workday.
 - ◆ Direct service operators to reduce power load to one of three levels according to a predetermined plan.
- ◆ Notification of pending/possible load shed will be made by service desk and computing liaisons.
- ◆ Time scale for putting load shed plan into operation:
 - ◆ 30 minutes from notification to providers on workdays.
 - ◆ 60 minutes from notification to providers on non-workdays if warning was given on previous workday.
 - ◆ 90 minutes from notification to providers on non-workdays if no warning was given.
- ◆ FOD will issue all-clear to resume full or partial operations when conditions permit.

What Will be Shutdown?

- ◆ Three possible load shed stages have been designated to lower power consumption by ~25%, 50%, and 99% depending on severity of the situation.
- ◆ Full racks will be shut down to reach desired power reduction starting with oldest, least efficient systems.
- ◆ There are three computer rooms at GCC:
 - ◆ GCCA: This room contains newer high density computing from 2012 purchases and will soon have the new LQCD nodes, but total power load in GCCA is relatively low. It is presumed that no actions will need to be taken in GCCA.
 - ◆ GCCB and GCCC contain older farm and LQCD nodes and both will be impacted by a load shed.
- ◆ Situation in GCCA will have to be re-evaluated as more nodes are added.

What Will Be Shut Down?

- ◆ Stage 1 25% load reduction target:

- ◆ GCCA: No action
- ◆ GCCB: Reduce percentage of cores available to clusters to:

CDF	CMS	D0	GP
71%	83%	76%	100%

- ◆ GCCC: Lower clock speed on nodes to attain required power reduction. If required reduction is not attained, then shut down designated nodes. Note this was adequate last summer to achieve ~25% load reduction.

- ◆ Stage 2 50% load reduction target

- ◆ GCCA: No action
- ◆ GCCB: Reduce percentage of cores available to clusters to:

CDF	CMS	D0	GP
55%	83%	68%	83%

- ◆ GCCC: Power down designated LQCD nodes to 50% capacity.

What Will Be Shut Down?

- ◆ Stage 3 99% load reduction target:
 - ◆ GCCA: No action
 - ◆ GCCB: Power down all racks except rack containing Fermigrid servers.
 - ◆ GCCC: power down all racks except 5 racks used by Lustre file system.
- ◆ FOD will monitor power consumption and may modify above plan based on the course of events during the shutdown.

Future Plans

- ◆ Basic problem is recirculating warm airflow around condensers on the ground next to GCC.
 - ◆ Soaker hoses helped in 2011, but that's just a bandaid.
 - ◆ Removing berm behind condensers also helped. Info from last summer indicates this lowered temps ~5 F. But temps still too close to danger zone.
 - ◆ Smoke tests last summer show building itself can significantly restrict air circulation when wind is in the right direction.
- ◆ Plan is to move condensers to roof of GCC where air flow is unobstructed. This will be done in non-disruptive manner.
- ◆ Executable plan expected to be ready by June.
- ◆ Goal is to have fixes in place by summer 2014 so computer rooms can be used at full capacity without load shed or non-standard mitigations.

For Additional Information

- ◆ More details may be found at:
 - ◆ DocDb: CS-doc-5015-v4
 - ◆ Contains full load shed plan and map of exactly what racks will be shut down at each stage.
 - ◆ DocDB: CS-doc-4793-v2
 - ◆ Contains detailed information regarding investigation into cooling problems at GCC and possible mitigations.
- ◆ Note: you will need a valid certificate in your browser to access these documents. You may also need explicit permission if you have never accessed Cdweb before. IF so, contact the Service desk to get access.